

I Sync, Therefore I DRAM

SILICON UPDATE

Tom Cantrell



These days, Intel is mainly known as the company that, in concert with Microsoft, controls what goes on inside most PCs. Their commanding position has served them well, propelling them to the top of the IC business.

After a decade of PC and x86 hoopla, it's easy to forget Intel's earlier contributions at the dawn of the silicon revolution. Long before the x86, Intel invented a chip that is arguably even more important—the DRAM.

Looking back at chips like the early '70s' Intel 1103 (see Figure 1a) is always good for a laugh. Imagine what your typical 1-MB PC would be like if it had 8000 of these little puppies. Forget the mini- and maxi-tower enclosures; the Trump tower is more like it. You'd better beef up your power supply too since each 1103 burned a half watt or so. Roll-your-own precharge, inverted DOUT, and 16-V I/O levels just add to the fun.

Nevertheless, as feeble as it now seems, the 1103 started the DRAM

ball rolling down a slippery slope of ever-increasing density and (at least until recently) ever-falling prices. Now, with the DRAM market moving through 4 and 16 Mb to beyond, a few chips are enough to handle the most bloated software upgrades without blinking a bit.

The modern DRAM era began in the late '70s with a 16-Kb DRAM such as the Intel 2116 illustrated in Figure 1b. More important than the technical features, acceptance of a standard kicked off multisourced DRAM as a commodity. The subsequent brutal battles for sockets had great benefits for users, but drove many suppliers (including Intel) from the business.

DRAM 101

Though the 2116 came (and went) before the x86, most of its features persist to this day. Thankfully, a single supply replaces the three (+5, +12, and -5 V) of yesteryear, but otherwise DRAMs haven't changed that much.

A typical 2ⁿ-bit DRAM is organized as 2^{n/2} rows and 2^{n/2} columns. As shown in Figure 2a, the row address is strobed into the DRAM with RAS* (row-address strobe). Next, the column address is strobed with CAS* (column-address strobe), which also clocks the data in or out depending on the WE* pin. The timing parameter t_{RAC} denotes the access time of the DRAM (i.e., t_{RAC} is 70 ns for a "70-ns DRAM").

To refresh your memory (ho-ho), DRAMs store each bit on a capacitor that leaks. Thus, they must be periodically refreshed, usually a row or two (limited by power considerations) at a time. Typical specs call for the entire

Figure 2a—Basic DRAM read and write cycles use a RAS*-before-CAS* technique. The t_{RAC} spec determines the access speed of the DRAM.


DRAM to be refreshed hundreds of times per second, though those pesky capacitors often prove surprisingly more bitworthy than that.

On average, a refresh cycle needs to occur every 15 μ s or so (i.e., 1-Mb DRAM needs 512 refresh cycles every 8 ms, a 4-Mb DRAM requires 1024 cycles every 16 ms, etc.). But, there's nothing that says you can't group all your refresh cycles in a burst.

Figure 2b shows the popular types of refresh cycles. The first is simply to read a location in each row, a scheme well-suited to graphics applications or software-refresh schemes. The similar RAS*-only refresh cycle was the most widely used on the 16-, 64-, and 256-Kb DRAMs. Starting with 1-Mb DRAMs, CAS*-before-RAS* (also known as CBR) refresh was offered and has since become popular. With an integrated refresh-address counter, CBR doesn't require an external address.

About the same time as the 1-Mb DRAMs came into existence, the formerly simple-minded DRAMs started to get brainy with the addition of high-speed access modes (shown in Figure 2c). All modes rely on the fact that once a row is accessed, it's possible to access different columns quickly, without requiring a new RAS* cycle.

Page mode repeats only the CAS* cycle for each new column address, cutting access and cycle time nearly in half. Note that the column address need not be sequential. Nybble mode latches four columns for subsequent access by CAS*. Since access is always sequential, no column address is required, making nybble mode even



faster than page mode. Static-column mode, as the name implies, offers speedy SRAM-like access to the column depending solely on the address, not CAS*.

Though a little slower than the others, page mode has proven the market favorite. Recently, a variant of page mode called Extended Data Out (EDO) has been introduced, which quickens page-mode cycles by leaving the output driver on between CAS* pulses. It looks like EDO is going to be widely adopted, further cementing page mode's position as the high-speed access mode of choice.

Subsequently, DRAM technology has migrated into a variety of specialized ICs including video

RAM, TV line memory, pseudo-SRAM (PSRAM), and so on. Nevertheless, the plain old, x1 or x4 page-mode DRAM reigns supreme in the marketplace.

Now, fueled by the insatiable demand for more bandwidth (for less \$, of course), a whole raft of clever new DRAMs are crawling out of the lab. One particular model appears headed for success—thanks not only to good features, but also to the warm, fuzzy glow multisource standardization imparts. Say "Hello" to the new kid on the block, the *synchronous DRAM*, otherwise known as SDRAM.

ROCK 'EM, CLOCK 'EM

JEDEC (Joint Electron Device Engineering Council), a formerly mild-mannered international standards committee, has emerged as a powerful force in the memory market.

Originally, they, like many standards organizations, reacted to market realities by documenting after-the-fact standards. However, as the demand for standardization has grown, their mediation and blessing has become very important. JEDEC played a major role in making commodity SDRAMs a reality by forging a peace agreement between headstrong contenders.




Figure 2b—The Read and RAS*-only refresh methods, which require an address, are being supplanted by CAS*-before-RAS*, which doesn't.





Figure 1—(a) The Intel 1-Kb 1103 DRAM was one of the first DRAMs. (b) The modern DRAM era started with 16-Kb DRAMs like the 2116.



Tom reminds us that in the beginning,

DRAM was created. And, after many days, Hitachi brought forth synchronous DRAM, which offers better timing, greater speed, and more bandwidth. Tom says that this is good.

Figure 2c—Of the three high-speed access modes, Fast-Page mode is most widely used.



Let's start by taking a look at one of the first SDRAMs out of the blocks, the 4-Mb, 66-MHz Hitachi HM5241605 (see Figure 3). Thanks to JEDEC standardization, most of the discussion applies to SDRAMs from other suppliers, though you've always got to be on the lookout for minor spec differences.

The first thing you'll notice is it's wide (x16), thankfully easing the frustration caused by the narrow x1 and x4 DRAMs. So far, it looks like SDRAMs (4 Mb and 16 Mb) will initially be offered in x4, x8, x9, x16, and x18 configurations. Later, I expect we'll see (64 Mb) x32.

SDRAMs reflect the modern downsizing trend in their use of TSOP (Thin Small Outline Package) surface-mount technology. Thanks to TSOP, the miniaturization possibilities are truly impressive—1/2 MB (2 MB for 16-Mb SDRAM) fits in a footprint about the same as the old DIP (=10 x 20 mm) and just a little thicker (1.2 mm) than a business card.

The easiest way to move bits faster with less power is to move them less far, so SDRAMs downsize the power supply to 3.3 V $\pm 10\%$. Now here's a place you need to check an individual SDRAM data sheet. It looks like some SDRAMs may achieve 5-V TTL compatibility by tweaking their AC specs. However, this capability may disappear in future-generation SDRAMs with

finer geometries. In any case, for the same speed and power reasons, whatever the SDRAM connects to is increasingly likely to be 3.3 V too.

Otherwise, the pin names certainly look familiar. There's RAS*, CAS*, WE*, CE*, and multiplexed addresses just like before. The data lines (I/O0-I/O15) are bidirectional, but the separate DIN and DOUT lines of the plain DRAMs were usually

connected together on the PCB anyway. The SDRAM features separate power for the memory array (V_{CC} and V_{SS}) and I/O pins (V_{CCQ} and V_{SSQ}) to help isolate noise.

DQMU and DQML are individual byte enables for the data bus, necessary to support an atomic, byte-write operation. Otherwise, a byte write would require a word read, byte mask, word write sequence.

It's the clock pin (CLK) and its enable (CKE) that are most important. After all, that's what the S in SDRAM is all about. While a plain DRAM is kind of passive, sitting around waiting for someone to tickle RAS* and CAS*, an SDRAM is a rather nervous critter that wants to blast some data on every clock edge (see Figure 4). With clock rates up to 66 MHz available now and with 100 MHz just around the corner, an SDRAM can handle the data rates demanded by superduper CPUs and gee-whiz graphics.




Figure 3—The Hitachi HM5241605 synchronous DRAM might be considered a New Age DRAM.

ELECTRONIC BANKING

Looking inside the SDRAM (Figure 5) reveals differences hidden behind the familiar-sounding pin names. The most obvious difference is that the memory array is split into two banks (selected using A9). As we'll see later, this supports a ping-pong, bank-switching scheme which maximizes bandwidth.

The other thing to note is that the familiar-sounding pin names—RAS*, CAS*, WE*, and so on—no longer actually connect to latches and buffers, but instead feed a control-logic block.

It turns out these pins don't actually control the memory transfer (remember, CLK does that), but should be interpreted as bits in an opcode as shown in Figure 6. The pins are sampled at each CLK cycle and inject a command into the SDRAM pipeline unless in a DESL (CS* high) or NOP (RAS*, CAS*, WE* high) state. The first command to get familiar with is MRS (Mode Register Set) since it defines the key SDRAM operating characteristics (see Figure 7).




Figure 4—Synchronous DRAMs rely on a high-speed clock (e.g., 66 or 100 MHz) to time all commands and data transfers.

Burst length refers to how many data transfers and clocks are associated with each READ or WRITE command, the choices being 1, 2, 4, 8, and full page (which, in this case, is 256). Shorter bursts end automatically, while the BST command terminates full-page

bursts. Burst type refers to the addressing order of the burst and is normally sequential, unless you're connected to a '486 or Pentium, which feature a unique(!) interleaved, cache line fill.

Write mode offers the option of burst read and write or burst read with

Find out how you can add intelligence to any home, at a cost that's within your budget.

LIVING WITH AN INTELLIGENT HOME will change the way you live.


Written by David Gaddis, author of *Understanding & Installing Home Systems*

This VHS video cassette retails for \$24.95 plus \$5 s&h.

It is being offered to Circuit Cellar INK subscribers for only \$17.95 plus \$4 s&h (U.S.)

ORDER TODAY!

Don't let this exciting technology opportunity pass you by!




Take a tour through modern technology in today's home...

Circuit Cellar, Inc.
4 Park St. • Vernon, CT 06066
Tel: (203) 8752751
Fax (203) 872-2204

Replace Four Conventional PC/104 Modules with One SuperXT™ CMF8680 cpuModule™

Embedded PC/XT Controller with
Intelligent Power Management



PC/104 Compliant
Actual Size: 5.6 x 3.8 x 0.5"

\$449
100 pcs.

- PC/XT compatibility with 286 emulation
- 14 MHz, 16-bit 8086 CPU
- +5V only; 1.6W at 14.3 MHz, 1W at 7.2 MHz
- Intelligent sleep modes, 0.1W in Suspend
- ROM-DOS and RTD enhanced BIOS
- Compatible with MS-DOS & real-time operating systems
- 1M bootable Solid State Disk & free software
- 4K-bit configuration EEPROM (2K for user)
- 2M on-board DRAM
- IDE & floppy interfaces
- CGA CRT/LCD controller
- Two RS-232 ports, one RS-485 port
- Parallel, XT keyboard & speaker ports
- Optional X-Y keypad scanning/PCMCIA interface
- Watchdog timer & real-time clock

Expand This Or Any PC/104 System
with the

CM106 Super VGA Controller utilityModule™

- Mono/color STN & TFT flat panel support
- Simultaneous CRT & LCD operation
- Resolution to 1024 x 768 pixels
- Displays up to 256 colors

\$223
100 pcs.

Speed Product Development with the
DS8680 Development System

Your DS8680 includes the CMF8680, CM102 keypad scanning/PCMCIA, CM104 with 1.8" 85MB hard drive, CM106 SVGA controller & DM5406 12-bit, 100 kHz dataModule™ in an enclosure with external power supply, 3.5" floppy, keyboard, keypad, TB50 terminal board, SIGNAL-VIEW™, SIGNAL-MATH™, MS-DOS, SSD software & rdLin™ for just **\$2950**.

For more information on our PC/104 and ISA bus products, call today.



Real Time Devices USA

200 Innovation Blvd. • P.O. Box 906
State College, PA 16803 USA
(814) 234-8087 / Fax: (814) 234-5218

RTD Europa • RTD Scandinavia

Real Time Devices is a founder of the PC/104 Consortium.

#131




Figure 5—Inside, synchronous DRAMs are split into two banks for maximum performance. Many of the control pins have familiar names, but now, they are little more than the opcodes for an on-chip logic that controls the actual memory access.

single write. The latter mode may be useful in graphics applications (e.g., a burst read refreshing the CRT and single writes coming from the drawing processor) or to support a CPU with write-through cache.

CAS* latency depends on how many clock cycles elapse between a request and transfer, with higher clock

rates demanding more delay. For instance, a particular-speed SDRAM may achieve 15-ns clock cycle and access times with CAS* latency 2, but only 30 ns with CAS* latency 1. Figure 8 summarizes the CAS* latency and burst-length options.

As shown, a transfer is initiated with an *ACTV* (activate) command,

Function	Symbol	CKE	n	A7	A6	A5	A4	A3	A2	A1	A0
Ignore command	DESL	H	X	H	X	X	X	X	X	X	X
No operation	NOP	H	X	L	H	H	H	X	X	X	X
Burst stop in full page	BST	H	X	L	H	H	L	X	X	X	X
Column address and read command	READ	H	X	L	H	L	H	V	L	V	V
Read with auto-precharge	READ A	H	X	L	H	L	H	V	H	V	V
Column address and write command	WRIT	H	X	L	H	L	L	V	L	V	V
Write with auto-precharge	WRIT A	H	X	L	H	L	L	V	H	V	V
Row address strobe and bank act	ACTV	H	X	L	L	H	H	V	V	V	V
Precharge select bank	PRE	H	X	L	L	H	L	V	L	V	X
Precharge all bank	PALL	H	X	L	L	L	H	X	X	X	X
Refresh	REF/SELF	H	V	L	L	L	H	X	X	X	X
Mode register set	MRS	H	X	L	L	L	L	L	L	L	V

Note: H: V_{HH}; L: V_{HL}; X: V_{HL} or V_{HL}; V: Valid address input

Figure 6—The CS*, RAS*, CAS* and WE* pins encode the synchronous DRAM instruction set.

A9	A8	A7	A6	A5	A4	A3	A2	A1	A0
OPCODE	0	LMODE	BT	BL					
A6	A5	A4	CAS Latency	A3	Burst Type	A2	A1	A0	Burst Length
0	0	0	R	0	Sequential	0	0	0	1 1
0	0	1	1	1	Interleave	0	0	1	2 2
0	1	0	2			0	1	0	4 4
0	1	1	3			0	1	1	8 8
1	X	X	R			1	0	0	R R
						1	0	1	R R
						1	1	0	R R
						1	1	1	F.P. R

F.P. = Full Page (256)
R is Reserved (inhibit)

Figure 7—The MRS (Mode Register Set) command defines synchronous DRAM basic operating characteristics such as burst length and CAS* latency.

which is pretty much like the RAS* cycle of old. In this case, it selects the bank (A9) and row (A0-A8). This is followed, after the requisite delay (t_{RCD}), by a *READ* or *WRIT* command that specifies the column address (A0-A7) as in yesterday's CAS*. After the previously specified CAS* latency, the data is blasted on each clock edge.

The *READ A*, *WRIT A*, *PRE*, and *PALL* commands all have to do with precharge. Since precharge has been automatic since the days of the 2116, it might seem like a step backward to make the system designer handle it.

Fortunately, if you'd rather not be bothered, the A versions of the *READ* and *WRIT* commands feature automatic precharge. However, as in plain DRAM, automatic precharge lengthens the basic cycle time. Fortunately, overt control of precharge in conjunction with the bank scheme does offer the opportunity for aggressive designers to hide precharge in one bank while transferring data from another.

That about covers the SDRAM instruction set, except for our old friend refresh, of which two variants are offered, depending on the state of the CKE pin. If CKE is high (i.e., the clock is enabled), an *AUTO REFRESH* mimics a plain DRAM's CBR-refresh cycle (i.e., an on-chip refresh counter means no external address is required).

SELF REFRESH (CKE low) is kind of a power-down mode, in which the

DRAM keeps itself refreshed while consuming amazingly little power (only 2 mA compared to its 50-100-mA active consumption). It's tempting to just leave the SDRAM in *SELF REFRESH* between accesses. Unfortunately, however, exit from *SELF*

REFRESH requires the issuance of 1024 cycles of *AUTO REFRESH* to update the refresh counter, so *SELF REFRESH* is only suitable for serious napping.

Figure 9, showing the allowed instruction sequences, looks complicated, but really isn't if you talk it through. When idle, the SDRAM should be refreshed 1024 cycles every 16 ms. An *ACTV* command (think of it as a RAS* cycle) selects a bank and row for action. Once activated, the row can be accessed freely, bursting data both ways from random column addresses. After all transfers are completed (or if the row address changes), exit is via a precharge cycle, taking the SDRAM back to the idle and refresh state.


Remember, the key point of the two-bank scheme is that each bank can be in a different state and the state transition intervals are faster between, rather than within, banks. For instance, the minimum *ACTV*-to-*ACTV* (i.e., cycle time) spec within a bank is 110 ns while between banks it's only 30 ns.

TURBO-128 THE NEXT GENERATION EMBEDDED CONTROLLER

NO DEVELOPMENT TOOLS REQUIRED

READY TO PROGRAM IN BASIC OR ASSEMBLY
Photronics Research introduces the T-128:
A True Single Board BASIC Development
System. The T-128 is based on Dallas
Semiconductor's new 8051-compatible DS80C320
With its 2K clock speed (25MHz) and 3K cycle efficiency, an
instruction can execute in 160ns: an 8051 equivalent speed
of 62.5MHz!!! Equally impressive is the T-128's high-speed NV-RAM
interface. Any of the 128K RAM may be programmed directly from a PC file through the console;
eliminating EPROMs and associated tools. Program Development has never been faster or more
convenient, even with the finest EPROM emulator. The T-128 features PORT 0 bias and EA-select for
DS87C320 upgrade.

- Dallas Semiconductor's DS80C320
- 300% more efficient than the 8051
- Three 16-bit Timers/Counters
- 13 Interrupts (8 Ext. 7 Int)
- A second 16-bit Data Pointer
- 384 Bytes of Internal RAM
- Programmable Watchdog
- Brownout Protection
- Power-Fail Reset/Interrupt
- Power-On Reset
- Fully supported by Franklin CS1



- Modified BASIC-52 Interpreter (BASIC-520)
- New Fast Enough for New Applications
- Stack BASIC Programs and Autorum
- CALL ASM Routines for Maximum Speed

- I/O
- Three 6-bit Parallel Ports
- Two Full-Duplex RS232 Serial Ports
- Decoded Device I/O Strobes
- 50-Pin Bus Connector

- UPGRADE
- DS87C320 processor (33MHz)
- Instruction cycle: 12ns
- 25 MIPS
- 8051 equivalent: 82.5 MHz
- Internal 16K ROM/1K SRAM

Comes Ready to Run
with power adapter/cable assembly.
Includes utility diskette with
DETAILED TECHNICAL MANUAL
\$199 in RTV.

109 Camille St. • Amite, LA 70422 • (504) 748-9911 • Tech Support (504) 748-7090 • FAX (504) 748-4242

#132




Figure 8—Though not shown on the figure, remember that higher CAS* latency offers a faster clock rate. Short-burst transfers automatically terminate, but full-page bursts must be explicitly terminated with the BST command.

Does your Big-Company marketing department come up with more ideas than the engineering department can cope with? Are you a small company that can't afford a full-time engineering staff for one-of-a-kind designs?

Steve Ciarcia and the Ciarcia Design Works staff may have the solution. We have a team of accomplished programmers and engineers ready to design products or solve tricky engineering problems. Whether you need an on-line solution for a unique problem, a product for a startup venture, or just experienced consulting, the Ciarcia Design Works is ready to work with you. Just fax me your problem and we'll be in touch.

REMEMBER...A CIARCIA DESIGN WORKS!
Fax (203) 871-8986

Unleash the Power of PromICE

- ❖ Connect via ROM Socket; DIP; PLCC; SMT
- ❖ Emulate ROMs up to 16MBit in Size
- ❖ Fastest Downloads Available: Parallel; Serial; Ethernet
- ❖ Run Industry Standard Debuggers
- ❖ Target Processor Independent
- ❖ Support 3Volt Targets
- ❖ Host Software Sources Included
- ❖ Shielded Cables for Reliable Operation
- ❖ 30-day Money-Back Guarantee; 1 Year Warranty!
- ❖ Unlimited Phone Support; 24hr BBS

Call Today 1-800-PROMICE (1-800-776-6423)

Grammar Engine Inc.
 921 Eastwind Dr., Suite 122 • Westerville, OH 43081
 614/899-7878 • Fax 614/899-7888

#133




Figure 9—An synchronous DRAM is either idle or actively reading and writing or precharging.

BURST OR BUST

You might think all this 66- and 100-MHz talk implies some kind of 10-times breakthrough in DRAM cell speed. In fact, the speed of basic memory operation really isn't much different (witness the 110-ns cycle time mentioned above is about the same as plain DRAMs). The SDRAM simply blazes a faster trail to what's otherwise regular DRAM.

Thanks to the ping-pong feature, even the simplest SDRAM setup (CAS* latency = 1, burst length = 1) offers good potential. Note that the parallelism can be extended across multiple SDRAMs as well (i.e., you can also ping pong between chips).

However, there's no doubt that bursty applications such as high-end, 32-bit CPUs and graphics best match the features of SDRAM. The tradeoff between CAS* latency, clock rate, burst length, number of chips, and so on is hard to generalize.

Graphics applications will likely shoot for maximum bandwidth (e.g., CAS* length = 3, burst length = full page, CLK = 66 MHz). It then becomes simply a matter of choosing the number of SDRAMs to support a given display. For example, a 1024 x 768 x 24 x 72-Hz display calls for 170 MB per second, while each 66-MHz SDRAM offers 132 MB per second. Thus, a two-SDRAM frame buffer easily handles

display refresh, leaving bandwidth for drawing. CPU main-memory are a little more tricky. A simple controller may just set the burst equal to the cache-block-refill size and leave it at that. CAS* latency versus CLK rate largely depends on memory-controller speed, cost and complexity issues, and "hit ratio" tradeoffs (i.e., weighing greater latency penalty against higher burst bandwidth).

More aggressive controller designs are possible if you're willing to sift through the entrails of your program and truly understand its behavior. It might be wise to put instructions and data in separate SDRAMs to take advantage of locality. Often, instruction- and data-cache-refill sizes are different, so each SDRAM's burst length could be set appropriately.

Another idea is to control burst length dynamically, either by overtly terminating full-page bursts or by reprogramming the mode register. This opens the door to really clever dynamic or "greedy" burst schemes, which wring all possible bandwidth out of the SDRAM.

No doubt, the prospect of rolling your own SDRAM controller, especially a fancy one, is rather daunting. The good news is I expect SDRAM support will quickly migrate into likely partners including 32-bit RISC CPUs, graphics ICs, and PC chipsets.

Tom Cantrell has been an engineer in Silicon Valley for more than ten years working on chip, board, and systems design and marketing. He can be reached at (510) 657-0264 or by fax at (510) 657-5441.

CONTACT

Hitachi America, Ltd.
 Semiconductor and IC Division
 2000 Sierra Point Parkway
 Brisbane, CA 94005-1819
 (415) 589-8300
 Fax: (415) 583-4207

IRS

422 Very Useful
 423 Moderately Useful
 424 Not Useful